

WHAT IS CLAIMED IS:

1 1. A multiprocessor computer system comprising a plurality of nodes,
2 each node including:
3 an interface to a local memory subsystem, the local memory subsystem storing
4 a multiplicity of memory lines of information and a directory, said directory configured to
5 store sharing information concerning a memory line of information stored in the local
6 memory subsystem;
7 a memory cache for caching a multiplicity of memory lines of information,
8 including memory lines of information stored in a remote memory subsystem that is local to
9 another node;
10 a protocol engine implementing a negative acknowledgment free cache
11 coherence protocol, the protocol engine including
12 a memory transaction array for storing an entry related to a memory
13 transaction, the entry including a memory transaction state, the memory transaction
14 concerning a memory line of information;
15 logic for processing the memory transaction, including advancing the
16 memory transaction when predefined criteria are satisfied and storing a state of the memory
17 transaction in the memory transaction array;
18 a first subset of nodes comprising one or more nodes from said plurality of nodes,
19 each node in said first subset including an interface to one or more I/O devices;
20 the protocol engine included in each node from said first subset of nodes is configured
21 to limit to a predefined period of time, any sharing of a memory line of information stored in
22 the remote memory subsystem that is local to a home node from said plurality of nodes; and
23 the protocol engine included in the home node is configured to identify only nodes
24 other than nodes in the first subset that are sharing the memory line of information.

1 2. The system of claim 1, wherein
2 the protocol engine included in the home node is further configured to maintain a
3 count of the nodes in the first subset that are sharing the memory line of information.

1 3. The system of claim 2, wherein

2 the protocol engine included in the home node is further configured to place a request
3 for exclusive ownership of the memory line of information in a pending state if the count
4 indicates that at least one node in the first subset is sharing the memory line of information.

1 4. The system of claim 3, wherein
2 the protocol engine included in the home node is further configured to remove from
3 the pending state the request for exclusive ownership of the memory line of information when
4 the count subsequently indicates that none of the nodes in the first subset are sharing the
5 memory line of information.

1 5. The system of claim 3, wherein
2 the protocol engine included in the home node is further configured to place in a
3 pending state a request to share the memory line of information from a requesting node of the
4 first subset of nodes if the request for the exclusive ownership is in a pending state.

1 6. The system of claim 5, wherein
2 the protocol engine included in the home node is further configured to remove from
3 the pending state the request to share the memory line of information from the requesting
4 node after the request for exclusive ownership of the memory line of information is removed
5 from the pending state.

1 7. The system of claim 2, wherein
2 the protocol engine included in the home node is further configured to increment the
3 count in response to a request to share the memory line of information from a requesting node
4 of the first subset of nodes.

1 8. The system of claim 2, wherein
2 the count is maintained in an entry of the directory, said entry being associated with
3 the memory line of information.

1 9. The system of claim 2, wherein

2 the protocol engine included in the home node is further configured to reduce the
3 count in response to a notice of invalidation from a requesting node of said first subset of
4 nodes, said notice of invalidation indicating that the requesting node is no longer sharing the
5 memory line of information.

1 10. The system of claim 1, wherein
2 the protocol engine included in each node in said first subset of nodes is configured to
3 store information in the memory transaction array to facilitate monitoring of the predefined
4 period of time.

1 11. The system of claim 10, wherein
2 modifying the memory transaction array includes encoding a time at which the
3 predefined period of time expires, said time subsequently being compared to a current time.

1 12. The system of claim 10, wherein
2 modifying the memory transaction array includes encoding a numerical value, said
3 numerical value subsequently being reduced a predefined number times such that the
4 predefined period of time expires approximately when the numerical value is reduced to a
5 predefined value.

1 13. The system of claim 10, wherein
2 the sharing of the memory line of information begins when a requesting node of said
3 first subset of nodes receives an affirmative response to a request to share the memory line of
4 information.

1 14. The system of claim 10, wherein
2 the protocol engine included in each node in said first subset of nodes is configured to
3 scan the memory transaction array to determine whether the predefined period of time has
4 expired.

1 15. The system of claim 14, wherein

2 the protocol engine included in each node of said first subset of nodes is configured to
3 terminate the sharing of the memory line of information if the predefined period of time has
4 expired.

1 16. The system of claim 15, wherein
2 terminating the sharing of the memory line of information includes invalidating a
3 copy of the memory line of information stored in the memory cache.

1 17. The system of claim 15, wherein
2 terminating the sharing of the memory line of information includes sending an
3 invalidation notice to the home node.

1 18. The system of claim 1, wherein
2 the protocol engine included in the home node is further configured to send an
3 invalidation acknowledgment to a node of the plurality of nodes requesting exclusive
4 ownership of the memory line of information after each node of the first subset of nodes that
5 was sharing the memory line of information terminates the sharing of the memory line of
6 information, said node from the plurality of nodes requesting exclusive ownership of the
7 memory line of information being prevented from completing a memory transaction
8 associated with the request for exclusive ownership until after each node of the first subset of
9 nodes that was sharing the memory line of information terminates the sharing of the memory
10 line of information.

1 19. The system of claim 1, where
2 the protocol engine included in the home node is further configured to limit the
3 sharing of the memory line of information by the one or more requesting nodes to a
4 predefined period of time.

1 20. The system of claim 19, wherein
2 the protocol engine included in the home node is further configured to set a time field
3 associated with the memory line of information in the memory transaction array to monitor
4 the predefined period of time.

1 21. The system of claim 19, wherein
2 the protocol engine included in each node in said first subset of nodes is configured to
3 scan the memory transaction array to determine whether the predefine period of time has
4 expired.

1 22. The system of claim 19, wherein
2 the protocol engine included in each node in said first subset of nodes is configured to
3 terminate the sharing of the memory line of information if the predefine period of time has
4 expired.

1 23. A multiprocessor computer system comprising a plurality of nodes,
2 each node including:
3 an interface to a local memory subsystem, the local memory subsystem storing
4 a multiplicity of memory lines of information and a directory, said directory configured to
5 store sharing information concerning a memory line of information stored in the local
6 memory subsystem;
7 a memory cache for caching a multiplicity of memory lines of information,
8 including memory lines of information stored in a remote memory subsystem that is local to
9 another node;
10 a protocol engine implementing a negative acknowledgment free cache
11 coherence protocol, the protocol engine including
12 a memory transaction array for storing an entry related to a memory
13 transaction, the entry including a memory transaction state, the memory transaction
14 concerning a memory line of information;
15 logic for processing the memory transaction, including advancing the
16 memory transaction when predefined criteria are satisfied and storing a state of the memory
17 transaction in the memory transaction array;
18 a first subset of nodes comprising one or more nodes from said plurality of nodes,
19 each node in said first subset including an interface to one or more I/O devices;

20 the protocol engine included in a requesting node from said first subset of nodes
21 configured to send to a home node from said plurality of nodes a request to share a memory
22 line of information stored in the remote memory subsystem that is local to the home node;

23 the protocol engine included in the home node further configured to generate a reply
24 in response to the request to share the memory line of information such that the requesting
25 node receives a copy of the memory line of information;

26 the protocol engine included in the home node further configured to update a count
27 maintained in a directory entry associated with the memory line of information in response to
28 the request to share the memory line of information, said count subsequently indicating that
29 an unidentified node from said first subset of nodes is sharing the memory line of
30 information;

31 the protocol engine included in the requesting node further configured to send a notice
32 of invalidation after a predetermined amount of time sharing the memory line of information;
33 and

34 the protocol engine included in the home node further configured to update said count
35 in response to the notice of invalidation, said count subsequently indicating that the
36 unidentified node from said first subset of nodes is not sharing the memory line of
37 information.

1 24. The system of claim 23, wherein

2 the protocol engine included in the requesting node is further configured to invalidate
3 the copy of the memory line of information after the predetermined amount of time sharing
4 the memory line of information.

1 25. The system of claim 23, wherein

2 the protocol engine included in the home node is further configured to place a request
3 for exclusive ownership of the memory line of information in a pending state if the count
4 indicates that at least one unidentified node is sharing the memory line of information.

1 26. The system of claim 25, wherein

the protocol engine included in the home node is further configured to remove from the pending state the request for exclusive ownership of the memory line of information after the count indicates that no unidentified nodes are sharing the memory line of information.

27. The system of claim 25, wherein

the protocol engine included in the home node is further configured to place a request by a node from the first subset of nodes to share the memory line of information in a pending state if the request for exclusive ownership of the memory line of information is in a pending state.

28. The system of claim 27, wherein

the protocol engine included in the home node is further configured to remove from the pending state the request by the node from the first subset of nodes to share the memory line of information after the request for exclusive ownership of the memory line of information is removed from the pending state.

29. A multiprocessor computer system comprising a plurality of nodes,
each node including:

an interface to a local memory subsystem, the local memory subsystem storing a multiplicity of memory lines of information and a directory, said directory configured to store sharing information concerning a memory line of information stored in the local memory subsystem;

a memory cache for caching a multiplicity of memory lines of information, including memory lines of information stored in a remote memory subsystem that is local to another node;

a protocol engine implementing a negative acknowledgment free cache coherence protocol, the protocol engine including

a memory transaction array for storing an entry related to a memory transaction, the entry including a memory transaction state, the memory transaction concerning a memory line of information;

15 logic for processing the memory transaction, including advancing the
16 memory transaction when predefined criteria are satisfied and storing a state of the memory
17 transaction in the memory transaction array;

18 a first subset of nodes comprising one or more nodes from said plurality of nodes,
19 each node in said first subset including an interface to one or more I/O devices;

20 the protocol engine included in a requesting node from said first subset of nodes is
21 configured to send to a home node from said plurality of nodes a request to share a memory
22 line of information stored in the remote memory subsystem that is local to the home node;

23 the protocol engine included in the home node further configured to generate a
24 response to the request to share the memory line of information such that the requesting node
25 receives a copy of the memory line of information, wherein the home node is configured to
26 retain sharer information identifying only nodes other than nodes in the first subset that are
27 sharing the memory line of information;

28 the protocol engine included in the home node further configured to set of field of a
29 directory entry associated with the memory line of information in response to the request to
30 share the memory line of information, said field subsequently indicating that an unidentified
31 node from said first subset of nodes is sharing the memory line of information;

32 the protocol engine included in the home node further configured to update the field
33 of the directory entry after a first predetermined period of time such that the field no longer
34 indicates that the unidentified node is sharing the memory line of information; and

35 the protocol engine included in the requesting node further configured to invalidate
36 the copy of the memory line of information after a second predetermined amount of time, the
37 second predetermined amount of time corresponding to the first predetermined amount of
38 time.

1 30. The system of claim 29, wherein

2 the protocol engine included in the home node is further configured to place a request
3 for exclusive ownership of the memory line of information in a pending state if the field of
4 the directory entry indicates that at least one unidentified node is sharing the memory line of
5 information.

1 31. The system of claim 30, wherein

2 the protocol engine included in the home node is further configured to remove from
3 the pending state the request for exclusive ownership of the memory line of information after
4 the field of the directory entry indicates that no unidentified nodes are sharing the memory
5 line of information.

1 32. The system of claim 31, wherein

2 the requesting node is a first requesting node; and

3 the protocol engine included in the home node is further configured to place a request
4 to share the memory line of information from a second requesting node from the first subset
5 of nodes in a pending state if the request for exclusive ownership of the memory line of
6 information is in a pending state.

1 33. The system of claim 32, wherein

2 the protocol engine included in the home node is further configured to remove from
3 the pending state the request to share the memory line of information by the second
4 requesting node after the request for exclusive ownership of the memory line of information
5 is removed from the pending state.

1 34. The system of claim 29, wherein

2 the protocol engine included in the home node is further configured to update the field
3 of the directory entry each time a request is received from a node from the first subset of
4 nodes to share the memory line of information while the field of the directory entry indicates
5 that at least one unidentified node is sharing the memory line of information.

1 35. A multiprocessor computer system comprising a plurality of nodes,

2 each node including:

3 an interface to a local memory subsystem, the local memory subsystem storing
4 a multiplicity of memory lines of information and a directory;

5 a memory cache for caching a multiplicity of memory lines of information,
6 including memory lines of information stored in a remote memory subsystem that is local to
7 another node;

8 a protocol engine implementing a negative acknowledgment free cache
9 coherence protocol, the protocol engine including
10 a memory transaction array for storing an entry related to a memory
11 transaction, the entry including a memory transaction state, the memory transaction
12 concerning a memory line of information; and
13 logic for processing the memory transaction, including advancing the
14 memory transaction when predefined criteria are satisfied and storing a state of the memory
15 transaction in the memory transaction array;
16 a first subset of nodes comprising one or more nodes from said plurality of nodes,
17 each node in said first subset including an interface to one or more I/O devices;
18 the protocol engine included in each node in said first subset of nodes configured to
19 treat requests initiated by said one or more I/O devices for a shared copy of a memory line of
20 information maintained in a memory subsystem that is local to another node as a request for
21 an exclusive copy of the memory line of information.

1 36. The system of claim 35, wherein
2 each node in said first subset of nodes does not include a processor core.

1 37. The system of claim 35, wherein
2 the protocol engine included in each node in said first subset of nodes is configured to
3 treat a read request as a read-exclusive request.